

Featured Articles

Collaborative Creation with Customers for Predictive Maintenance Solutions on Hitachi IoT Platform

Chetan Gupta, Ph.D.
 Ahmed Farahat, Ph.D.
 Tomoaki Hiruta
 Kosta Ristovski, Ph.D.
 Umeshwar Dayal, Ph.D.

OVERVIEW: Through the proliferation of sensors, smart machines, and instrumentation, industrial operations are generating ever increasing volumes of data of many different types and our customers are demanding solutions that provide business value over this collected data. In our interactions with customers across verticals, we have discovered that there is an urgent need for predictive maintenance solutions that meet customer demands. The reason for the appeal of predictive maintenance solutions is their ability to increase equipment availability, reduce the cost of unexpected failures and make operations more predictable. Hitachi offers a portfolio of data analytics technologies to address predictive maintenance use cases in a variety of verticals and in this paper we present an overview of our work in this area.

INTRODUCTION

WE are at the cusp of transformative changes across industries, from agriculture to manufacturing, from mining to energy production, from healthcare to transportation. These transformations hold the promise of making our economic production more efficient, cost-effective, and sustainable and are being driven by the convergence of the global industrial system operational technology (OT) with the power of integrating advanced computing, analytics, low-cost sensing and new levels of connectivity information technology (IT). This convergence enables the creation of a new class of big data solutions for monitoring, managing, and optimizing industrial operations and physical systems.

At Hitachi, we realize that there is a need for such solutions in the broader marketplace. We can use our decades of experience in equipment manufacturing and marry it with our expertise in analytics to bring unique solutions to market that solve some hard and important customer problems. The Global Center for Social Innovation (CSI) has been chartered to accomplish this vision. Within CSI, we have developed a methodology for creating new solutions that are of value to customers (see Fig. 1). It starts with having a dialogue with customers to understand their needs and pain-points using a design methodology. The next step is the development of prototype solutions and demos

followed by a proof-of-value at the customer site. To build these prototype solutions, we use both in-house CSI technology as well as technology from the Center for Technology Innovation (CTI). The final step is working with a business unit to scale the solution to multiple customers in the same vertical as well as to customers across different verticals.

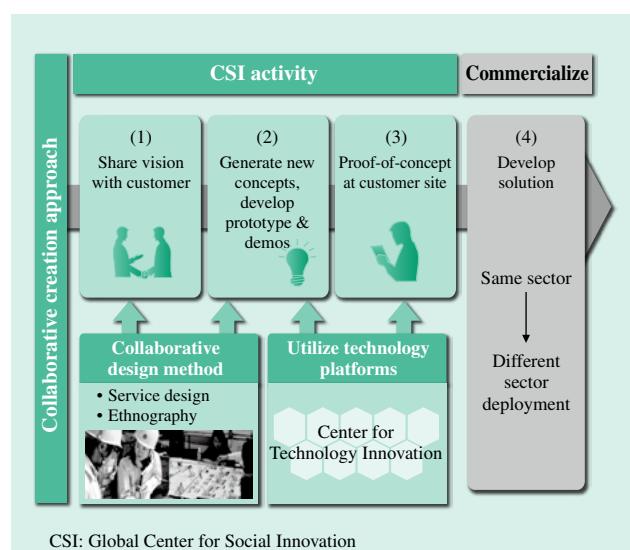


Fig. 1—Collaborative Creation with Customers for Predictive Maintenance Solutions.

The diagram shows the stages of the collaborative creation with customers methodology used to create new solutions that are of value to customers.

One application area that is common across verticals is the area of predictive maintenance – the ability to do the “right maintenance at the right time.” This includes problems such as performance monitoring and modeling, maintenance analytics, and of course, failure prediction. Predictive maintenance allows operators to reduce downtime and the cost of unexpected failures, and increase availability, predictability, and confidence in their operations. For equipment vendors, it opens avenues for new service and business models.

This article presents an overview of the predictive maintenance technologies created by Hitachi in close collaboration with customers. The article demonstrates practical use cases and describes predictive maintenance solutions that will be available on the Hitachi Internet of Things (IoT) platform, called Lumada, and a common analytics framework.

The rest of the article is organized as follows. The section entitled “Data Analytics for Predictive Maintenance,” gives an overview of predictive maintenance and the solutions offered by Hitachi to its customers. “Failure Prediction Use Cases” focuses on use cases for failure prediction, which is a key problem in predictive maintenance. “Solutions on Common Analytics Framework and Lumada” describes how to build repeatable achieve predictive maintenance solutions. And “Conclusions” concludes the paper.

DATA ANALYTICS FOR PREDICTIVE MAINTENANCE

Maintenance is a process for which the objective is to keep the equipment in a working, efficient and cost-effective condition. The maintenance process is conducted by performing the necessary actions on the equipment to achieve one or more of these objectives. These actions include, but are not limited to, inspection, tuning, repair, and overhaul of the equipment or its components.

Predictive maintenance is a maintenance strategy that depends on monitoring the condition of the equipment in order to determine the right maintenance actions that need to be taken, at the right time. Predictive maintenance has many advantages over other strategies such as corrective and preventive maintenance as it reduces the chance of unexpected failures, increases the equipment availability, and accordingly, decreases the overall cost of the maintenance process.

Predictive maintenance technologies utilize one or both of the following.

(1) Physical devices⁽¹⁾

Physical devices that assist in the diagnosis of equipment conditions such as devices traditionally used for vibration monitoring, lube oil analysis, particle wear analysis, thermography, and ultrasonic analysis.

(2) Software technologies

Software technologies that continuously monitor and analyze the sensor and event data generated by the equipment and other physical devices along with maintenance and operation data.

Software technologies for predictive maintenance can be further classified into:

(a) Knowledge-driven systems

In these systems, information about conditions to be monitored (e.g., pre-failure conditions) are manually encoded by the equipment manufacturer or other experts in the equipment domain. This is the most common technology for software-based predictive maintenance that is usually embedded by the manufacturer into the control software of the equipment. This technology is however limited by the knowledge of domain experts about possible patterns of equipment conditions or by the sophistication of the physical models used at design time. For complex equipment, such methods fail to capture the interrelationships between the sub-components, and as a result, the actual behavior of the equipment is often different from physics based simulation models or pre-defined human generated rules.

(b) Data-driven systems

In these systems, information about the conditions of interest are learned from historical sensor data and event logs. In comparison to knowledge-based technologies, data analytics have the potential of capturing complex patterns that are often not captured by domain experts during the design or deployment time.

Data analytics technologies for predictive maintenance can be categorized into descriptive, predictive, and prescriptive. The rest of this section describes some of the technologies offered by Hitachi in each of these categories.

Descriptive Analytics

Descriptive analytics discover actionable insights from historical data about the operation of the equipment. These insights enable the maintenance personnel and management to improve the maintenance process, and eliminate inefficiencies in the equipment operation. Examples of descriptive analytics technologies for predictive maintenance are described below.

Performance degradation detection

We initially developed this technology for a customer who wanted to understand the performance degradation of a set of equipment that had been in the field for close to 10 years. This involves detecting slow degradation in the performance of the equipment or its components. This slow degradation could be an early warning of some failure, or could reflect an inefficient or cost-ineffective state of the equipment that needs to be addressed. Our solution for performance degradation detection learns the ideal performance of the equipment from historical data based on predefined health indices. Health indices are usually defined in collaboration with domain scientists and then normalized using machine learning in order to remove the effect of load and seasonality. By continuous monitoring of normalized health indices, it is possible to detect performance degradation. This results in an early enough warning for the maintenance staff to take actions to prevent a failure or restore the equipment to an efficient and cost-effective state.

Maintenance effectiveness estimation

Our technology for maintenance effectiveness estimation depends on statistical analysis of the performance before and after maintenance actions⁽²⁾ and was initially developed in close collaboration with a customer who had regular maintenance schedules for their equipment, and wanted to understand if they were really being effective with their practices or not. This analysis determines whether an individual or a class of maintenance actions resulted in statistically significant improvement of performance or not. This sort of analysis is very valuable to the maintenance staff. For example, the maintenance staff can get feedback about the past and ongoing maintenance actions and learn whether they were/are successful in keeping the equipment in the desired condition. By knowing that a particular maintenance action did not improve the performance of some component of the equipment as it was supposed to, the maintenance staff can quickly implement an effective countermeasure. Furthermore, these insights, based on actual measurements from the equipment, can help maintenance operators and their management to improve day-to-day maintenance actions and revise the maintenance process for a fleet of equipment (e.g. change the provider).

Predictive Analytics

These technologies are mainly concerned with the prediction of future events such as failures, based on learning over historical data. Failure prediction

is a key problem in predictive maintenance which belongs to the category of predictive analytics. We offer an extensive portfolio of algorithms for failure prediction that handle a variety of use cases. A detailed description of these technologies is the subject of “Failure Prediction Use Cases.”

Prescriptive Analytics

These technologies generate recommendations for the maintenance personnel or management that result in reduction of failure rates while meeting operational objectives. Examples of prescriptive analytics include the following.

Operating envelope recommendation

This technology learns, from the historical data, the subset of operating conditions that results in the reduction of failure rate while achieving the operational targets. (Similar approaches can be used to reduce the cost of operations too, but they do not fall under the umbrella of predictive maintenance). For instance, we have used correlation analysis to study the effect of operating conditions of heavy-duty trucks on their failure rates for one of our customers who was experiencing higher than usual failure rates and wanted to reduce the failure rate while maintaining its existing production. These correlations are then used to construct rules that provide recommendations to truck drivers on specific operational conditions such that the failure rate will decrease in the future.

Maintenance optimization

As both analytics and human-based predictive insights, predictions, and recommendations increase, it is sometimes difficult for the maintenance team to know which maintenance action to take or to understand the global impact of a particular maintenance action. For instance, two predictors might recommend immediate maintenance of two distant pieces of equipment at the same time, which might be impossible given limited maintenance resources. To solve this problem, we are developing an overall optimization framework that takes into consideration the outputs of the predictive maintenance algorithms along with different cost estimates and operation constraints, and recommends a complete maintenance plan that results in maximum operation efficiency and minimum maintenance cost.

FAILURE PREDICTION USE CASES

Failure prediction is a key problem in predictive maintenance, and is concerned with estimating the likelihood that an undesirable condition or event is

going to happen early enough so that a countermeasure can be implemented to prevent the condition/event from happening. Failure prediction techniques mainly depend on encoding information about pre-failure conditions of the equipment and then monitoring real time sensor and event data searching for these conditions. We offer a portfolio of techniques for failure prediction that address a variety of use cases. This portfolio includes event-based, sensor-based, and model-based failure prediction.

Event-based Failure Prediction

In cases when collection of raw sensor data is expensive or infeasible, usually only event data is available for analysis. In these cases, time-stamped events are generated based on real-time sensor data at the equipment level and then transmitted to the operational database. We have developed technology that treats these events as long temporal sequences that can be mined for possible relationships using association rule mining⁽³⁾.

Although the initial motivation behind association rule mining was to analyze market basket data, new approaches have been developed to address problems in various domains such as environmental monitoring, bioinformatics, telecommunications, etc. These approaches tackle the problems such as mining frequent or rare patterns in temporal or non-temporal sequences of events. Our technology modifies one of the existing approaches and mines the temporal event sequences to find significant co-occurrences of events

TABLE 1. Failure Prediction Rules

Rules for failure prediction that were learned with heavy mobile equipment.

Event	Description	Prediction	Impact	Confidence
E71, E72, E83	Electrical system, Engine, Tires	S23	Standby	99%
E71_A	Engine	S23	Standby	96%
E77	Hydraulic oil leak	S23	Standby	95%
E71, E78	Electrical system, Propulsion	E72	Engine	64%

within pre-specified time windows. Once identified, these co-occurrences will define temporal association rules that can predict future failures. Then, these rules can be applied to the incoming events in order to predict the potential occurrence of failures within a time-window that makes business sense, namely it allows for corrective maintenance actions. The technology deployment pipeline is presented in Fig. 2.

We have applied our technology on event data generated by heavy mobile mining equipment for one of customers and some rules with the highest confidence from the analysis are shown in Table 1. These rules often provide significant value. For example, even in the absence of controlled area network (CAN) bus data from vehicles' sensors, we can successfully predict engine problems based on the fourth rule (see Table 1). This rule indicates that if some equipment experiences electrical system and propulsion problems, then it will have engine problems within 10 days with 64% confidence (the pre-specified window length was 10 days).

Our results on heavy mobile equipment demonstrate the potential of our technology for failure prediction use cases when only equipment event data is available.

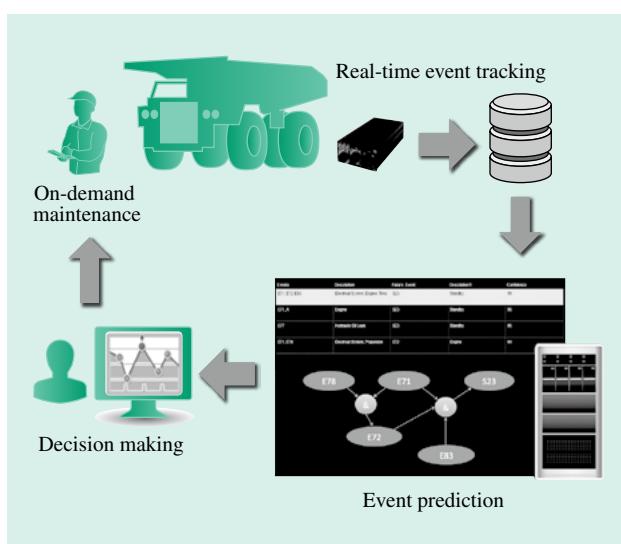


Fig. 2—Event-based Failure Prediction System.
Association rule mining over a temporal sequence is used to learn prediction rules for failures.

Sensor-based Failure Prediction

Sensor measurements, when available, encode rich information about pre-failure conditions. First, using machine-learning algorithms, failure prediction models are generated based on the pre-failure conditions learned from historical sensor measurements. Next, these models are applied to real-time data in order to predict failures. Two approaches to failure prediction from sensor data are given below:

Failure prediction using anomaly detection

In this prediction method, models of the normal behavior of a group of sensor measurements are learned from historical sensor data, and then any deviations from this normal behavior are detected during equipment operation. The anomaly detection algorithm learns normal clusters of data based on

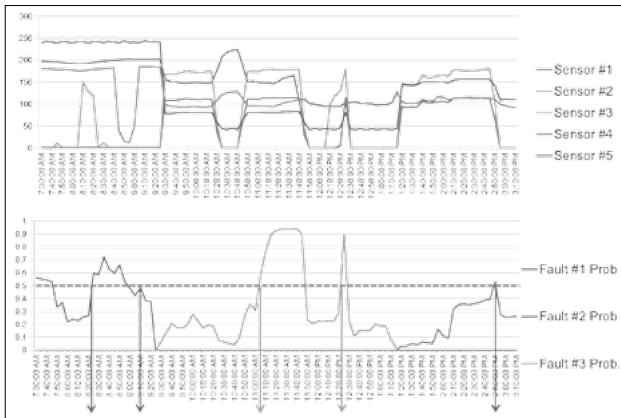


Fig. 3—Sensor-based Failure Prediction.

Sensor measurements are provided to multiple classifiers which calculate the probability that different failures will happen in the future.

the spacing between sensor measurements, and then calculates an anomaly score for each new measurement based on the distance of this measurement from the pre-learned clusters. This algorithm is suitable for use cases where there are a lot of sensor measurements related to normal operation, but not enough historical incidents of failures. Hitachi has successfully applied the algorithm to predict failures for a variety of customer use cases including heavy equipment such as generators.

Failure prediction using classification

In this prediction method, classification algorithms⁽⁴⁾ are used to learn complex pre-failure patterns based on historical failure incidents. The algorithm partitions the historical time series data into normal and pre-failure windows, and then derives a binomial classifier that differentiates between normal and pre-failure conditions. In comparison to the anomaly detection method, this method is able to recognize the particular type of failure, but it needs to have a sufficient number of samples of historical incidents for each failure type. This method has been quite successful in Hitachi customer environments. It has used this method, among other things, to predict failures in thermodynamic equipment such as chiller systems, and in vehicles maintained by equipment vendors and fleets. Fig. 3 shows a subset of sensor data used for predicting failures in thermodynamic equipment along with the probability estimated by the predictor.

Model-based Failure Prediction

Classification-based failure prediction models require a number of failure instances of different types from historical data. In cases where failures do not occur

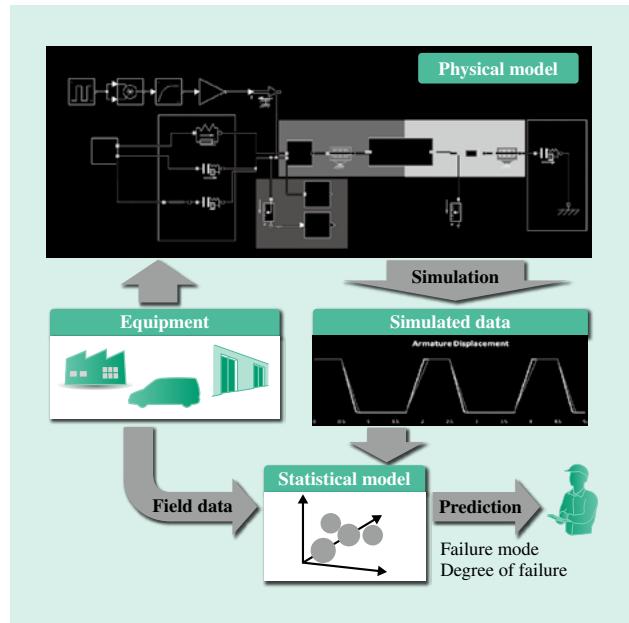


Fig. 4—Physical model-based failure prediction.

The physical models of equipment generate simulated data regarding normal and faulty behavior.

frequently in the field, Hitachi's technology leverages physical models of equipment to simulate these failures (see Fig. 4).

The physical models are defined using mathematical equations that represent the electrical and mechanical systems of the equipment. These models are built through collaboration with internal domain experts and through open innovation with partners. One of Hitachi's partners has been PARC (Palo Alto Research Center), a Xerox Company, with whom it has jointly developed physical models of various system components. Moreover, Hitachi has included some failure modes and the degree of failure in the physical models. Therefore, its physical models can simulate data corresponding to pre-specified failure modes and pre-specified degrees of failure.

Having physical models, it is possible to augment field-generated data with simulated data to learn more accurate classification-based predictive models. Hitachi has successfully utilized its domain knowledge to create physical models along with data analytics to improve predictive maintenance algorithms in several domains.

SOLUTIONS ON COMMON ANALYTICS FRAMEWORK AND LUMADA

As we have seen, there is a need for different predictive maintenance solutions across verticals. The traditional approach to meeting such a need is to create bespoke

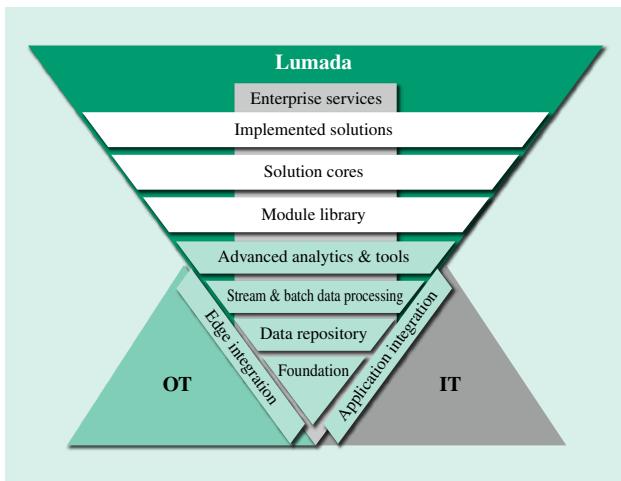


Fig. 5—Lumada IoT Platform.

The Hitachi IoT platform is used for building predictive maintenance solutions.

solutions from problem to problem and industry to industry. In contrast to this approach, we are building a common framework⁽⁵⁾ and integrating into Hitachi's IoT solution platform called Lumada⁽⁶⁾ (see Fig. 5) to achieve a predictive maintenance framework. This will allow for the creation of repeatable predictive maintenance solutions. Our predictive maintenance framework is geared towards processing sensor and event data, and provides a rich library of predictive maintenance solution components. Our approach of having a predictive maintenance framework offers us several advantages:

(1) For predictive maintenance solutions, the underlying architecture is often similar across verticals. For example, solutions need to process sensor or event data (or both), maintenance data, etc. Our framework provides the ability to manage, process, and analyze different types of data from structured to unstructured data, and from stored to streaming data. Specifically, it brings together data integration, advanced analytics, and visualization capabilities tailored to the predictive maintenance use cases. This dramatically reduces the time, effort, and money to build a new predictive maintenance solution.

(2) It allows us to utilize knowledge gained in one domain for another domain through reusable analytics solution components. For example, we discovered that a performance degradation technology that we developed for chillers could also be used for automobiles, simply by changing the KPI. Once we encoded the performance degradation methodology developed for the chiller use case as a solution component, to develop the performance degradation

solution for the automobile industry, we could simply invoke the previously-built solution component.

(3) Analytics applications in the Industrial space often have limited applicability and adoption because they are designed by and for IT experts. Our framework for predictive maintenance is designed to empower domain experts. It does this by allowing the domain experts to express their analytics intuition and build their own dashboards with a very simple user interface.

CONCLUSIONS

In this paper we presented a portfolio of data analytics technologies for predictive maintenance that were developed in close collaboration with customers.

Our technologies use data mining and machine learning algorithms to discover actionable insights about the history of the equipment, predict failures before they happen, and recommend countermeasures to prevent these failures from happening again. Our solutions are built on top of Lumada and a common framework that offers a library of predictive maintenance solution components.

REFERENCES

- (1) R. K. Mobley, "An Introduction to Predictive Maintenance," Second Edition, Butterworth-Heinemann (2002).
- (2) R. L. Ott et al., "An Introduction to Statistical Methods and Data Analysis," Cengage Learning (2010).
- (3) P.-N. Tan et al., "Introduction to Data Mining," First Edition, Pearson, pp. 327-390 (2005).
- (4) T. Hastie et al., "The Elements of Statistical Learning: Data Mining, Inference, and Prediction," Springer Series in Statistics (2009).
- (5) C. Gupta et al., "Analytics-Driven Industrial Big Data Applications", Japan Industrial Management Association, Management Systems **25**, No. 2.
- (6) Hitachi Press Release, Hitachi Unveils Lumada Internet of Things Core Platform (May 2016), <https://www.hitachiinsightgroup.com/en-us/pdf/press-release/hitachi-unveils-lumada-internet-of-things-core-platform.pdf>

ABOUT THE AUTHORS

**Chetan Gupta, Ph.D.**

*Big Data Lab, Global Center for Social Innovation
North America, R&D Division, Hitachi America, Ltd.
He is currently the Chief Data Scientist and
Department Manager at the Big Data Lab.*

**Ahmed Farahat, Ph.D.**

*Big Data Lab, Global Center for Social Innovation
North America, R&D Division, Hitachi America, Ltd.
He is currently engaged in the development of new
technologies for predictive maintenance.*

**Tomoaki Hiruta**

*Big Data Lab, Global Center for Social Innovation
North America, R&D Division, Hitachi America, Ltd.
He is currently engaged in the development of new
technologies related to predictive maintenance.
Mr. Hiruta is a member of The Japan Society of
Mechanical Engineers (JSME).*

**Kosta Ristovski, Ph.D.**

*Big Data Lab, Global Center for Social Innovation
North America, R&D Division, Hitachi America, Ltd.
He is currently engaged in the development of
new technologies for fleet analytics and predictive
maintenance.*

**Umeshwar Dayal, Ph.D.**

*Big Data Lab, Global Center for Social Innovation
North America, R&D Division, Hitachi America, Ltd.
He is currently a Senior Fellow and manages the
Silicon Valley Research Center, Global Center for
Social Innovation North America.*